Common Sense Investing: Bridging the Gap Between Expert and Novice

Ashwani Kumar
MIT Media Laboratory
Cambridge, MA
ashwani@mit.edu
617-253-3716

Sharad C. Sundararajan
IBM, Design Automation
Fishkill, New York
sharads@us.ibm.com
781-899-0341

Henry Lieberman
MIT Media Laboratory
Cambridge, MA
lieber@media.mit.edu
617-253-0315

ABSTRACT

In this paper, we describe Common Sense Investing (CSI), an interactive investment tool that uses a knowledge base of common sense statements in conjunction with domain knowledge to assist personal investors with their financial decisions, primarily asset-allocation. In interfaces that provide expert advice, one key problem is *elicitation* – how to ask questions that enable the expert model to make decisions, and at the same time, are understandable to the novice. The second problem is *explanation* – how to explain rationale behind expert decisions in terms that the user can understand. Many programs already encode expert models, but few have good models of novice knowledge, especially where broad knowledge of everyday life might bear on the subject. OMCSNet[1], a semantic network representation of the OpenMind Common Sense Knowledge Base[2], is the source of a wide range of facts about day-to-day life. CSI maps the user's goals, expressed in concepts from OMCSNet, to the expert's goals, expressed in technical financial terms. Instead of asking "What is your tolerance for risk?" where the user might not understand the concept of risk tolerance, we can ask, "Do you usually have a lot of credit card debt?" Aligning the expert's questions and decisions with common sense knowledge pertinent to the user increases the user's confidence in the ability of the system to meet their needs.

Keywords

CSI, Common Sense Reasoning, Investment, Frame Semantics, Information filter

ACM Classification Keywords

H.5.2. User Interfaces. H.5.3. Web-based interaction

INTRODUCTION

There exist numerous investment tools [1,2,3,4,5,6,7] that claim to come up with the best strategies for asset allocation through a sequence of questions to gauge the user's inclination towards investment and willingness to take risks. However, these tools suffer from various shortcomings such as lack of control to the user, limited personalization, and insufficient explanation of the rationale behind the decisions. They have complex models of the expert knowledge but eventually, fail in disseminating this knowledge in any useful manner to the novice. Effectively, there is very little effort in trying to bridge the gap between expert and novice. Common Sense knowledge about the expert domain and the user's day-to-day life can be used to overcome these shortcomings.

Common sense, as we understand it in today's world is the shared knowledge that establishes common ground. Common sense can be used to check the validity of various task scenarios and to help troubleshoot problems along the lines of Woodstein [8], an interactive E-Commerce debugger. There is a huge effort going on in this direction at the MIT Media Lab, to build a database with millions of common-sense facts that people come across in everyday life .

In this paper we suggest a unique approach to bridge the gap between novice and expert knowledge systems by providing an intuitive interactive framework where the user can interact with the system using natural language sentences without being overwhelmed by the expert knowledge processing that the system performs. In the following section, we describe the functional interface, followed by the system description with brief descriptions of various components in subsequent sections.

COMMON SENSE INTERFACE

Most of the existing investment advisory tools use inflexible interfaces like menus, drop-downs, buttons etc. providing only a handful of options that are sometimes not immediately clear or intuitive and may potentially collect superfluous information. Instead of asking the user how risk tolerant (s)he is by offering some discrete and incomprehensible levels of risk (Figure 1), it will be enormously beneficial to provide some explanation about the ramifications of being aggressive or taking high risks [9]. Common Sense facts like "frequent gambling means high risk" and "having recurring credit card debts means high risk" capture human tendencies of risk taking behaviors.

CHI 2004, April 24–29, 2004, Vienna, Austria.

ACM 1-58113-703-6/04/0004.

1 http://web.media.mit.edu/~hugo/omcsnet/index.html

2 http://commonsense.media.mit.edu/cgi-bin/search.cgi

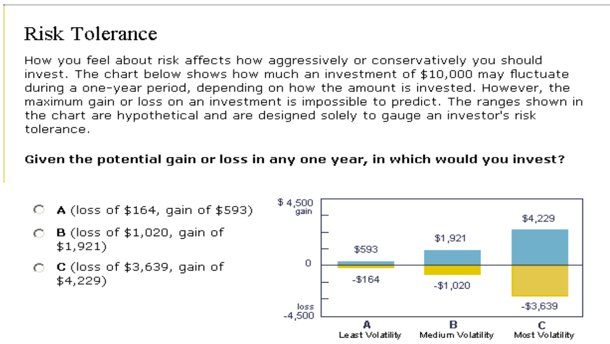

Figure 1: Vanguard's Risk Tolerance Question

In the following is an illustration of a typical interaction between the user and CSI, where CSI uses common sense knowledge to ask relevant questions and provide explanations (Figure 2):

User : I have 10,000 dollars.

CSI : Do you want to invest the full amount?

User : Yes

CSI : Are you in any kind of debt?

[CSI explains the question about debt as "Debts may have very high interest rates and it might be worth paying attention to these before using the entire amount for investment"]

User : No

CSI : How would you like to use this investment?

[This question is to figure out the user's end goal]

User : I want to buy a house.

CSI : How much will it cost?

User : $300,000.00

CSI : What is your timeline?

User : 5 years

[As the calculated growth rate turns out to be much more than the industry average, CSI qualifies it as "aggressive". Further, CSI uses the expert common sense knowledge to map "aggressive" to "risky" and asks:]

CSI : This seems aggressive. Do you tend to incur a lot of credit card debts?

Based on the reply, CSI can gauge what level of risk the user is willing to take and accordingly map it to the expert knowledge.

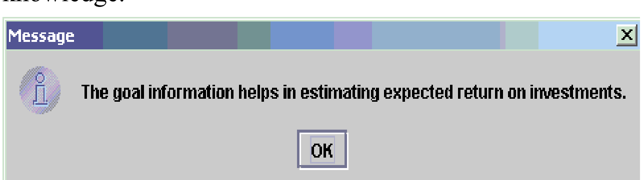

Figure 2: A Sample Explanation Message

SYSTEM DESCRIPTION

The system architecture consists of the following components:

1. Common Sense Analyzer (CSA)
2. Expert financial engine
3. Information Filter (InfoFilter)
4. OMCS Interface
5. Investment Strategist

CSA comprises of a Natural Language Understanding front-end, which processes the user's commands in natural language to build investment and goal structures [10]. The CSA implements an interface, which interacts with OMCSNet using SOAP[3] (Simple Object Access Protocol) messages. OMCSNet is configured to run as a web service, which is queried to extract semantic associations in the form of predicates. The investment strategist interprets the analyzed request from the CSA and accumulates relevant information from the InfoFilter and Expert Financial Engine.

Common Sense Analyzer

Contemporary research in the area of interactive systems has emphasized the importance of having a dialogue-based interaction as opposed to fixed menu or scenario based interactions with the user [11]. However, having dialogues in the mode of natural languages requires that the system have adequate language understanding capabilities, fail-soft inference and deduction mechanisms.

The key idea is to specify appropriate mappings from natural language utterances to expert system behaviors and vice versa. The important thing to note here is that these mappings are dynamic in the sense that they evolve with interactions, are personalized based on the user's profile, and get refined as the common sense knowledge base gets richer.

Natural Language Understanding (NLU) Unit

All the user's requests are first tagged using a Parts of Speech (POS) Tagger[4]. The tagged text is chunked using a text-chunker, which groups tagged words within an utterance to disjoint classes based on pre-defined rules. Further, a semantic analyzer produces the semantic parse of the sentence in the form of an n-ary argument structure (Figure 3).

```
---------------Tagging User Request---------------
I/PRP want/VBP to/TO invest/VB 1000/CD dollars/NNS

-------------- Chunking User Request ----------
(NX I/PRP NX) (VX want/VBP to/TO invest/VB VX) (NX 1000/CD dollars/NNS NX)

-------------- Semantic Parse of the request in the form: (Verb-Subj-Obj-Obj) -----------
("invest" "I" "1000 dollar")
```

Figure 3: Natural Language Parse

Based on the verbs occurring in the semantic parse and corresponding synonyms, the NLU unit constructs a frame-based semantic structure [12,13,14,15], which is then correlated with the lexical predicates in the OMCSNet. The frame structure comprises of hierarchical event-object structures derived from the semantic parse and chunked-text. This kind of generic type-based construction has

---

3 http://www.w3.org/TR/soap/

4 http://web.media.mit.edu/~hugo/montylingua

subsequent positive implications on goal planning and iterative interaction with the OMCSNet [16].

Action Planning

The system needs to map the derived semantics from the user's utterance to the intentional goal structures. As the investment-strategy process is iterative and complex, it is modeled using finite state automata, where states are characterized by the various steps needed to lay out an investment strategy and the transitions encode various choices that the user can express using natural language. Essentially, goals have slot-filler type structures and by progressing through the state automata, it is made feasible to attain the level where adequate investment advice could be extracted from the expert system. For instance, the frame structure for the "invest_goal" looks as following (Figure 4).

```
<<<Frame Name: invest>>>
        Type: event
        Subject: 1
        Objects:
        Object 1: <<obj1>>
                Type: dollar
                Attributes:
                Attribute 1:<<attr1>>
                Attribute_Name: individuation
                Attribute_Value: 1000
```

Figure 4: Frame Representation for <invest>

Similarly, the invest_action requires an "invest" frame, where the slots pertaining to the investment object is filled with the money to be invested. Naturally, maintaining frame semantics of an utterance has advantages as the utterance frames can be compositionally correlated with the action frames (for example, subsumption criteria).

Common Sense Inference

Frame semantics is an elegant framework for characterizing fully specified semantics. However, due to the inherent ambiguity and potential for multiple senses, it becomes essential to correlate the fully specified frame semantics to the relevant senses. Also, from the action planning point of view it is necessary to articulate necessary and sufficient steps to achieve the desired goal (Figure 5). The common sense knowledge fulfills both of these requirements as it encodes multiple senses in a semantic network, where traversal along a particular path provides various steps needed to complete a particular goal [17,18,19].

```
<<<Frame Name: buy>>>
        Type : event
        Subject : USER
        Objects:
        Object 1 :<<obj1>>
                Type : THING
                Attributes :
                Attribute 1:<<attr1>>
                Attribute_Name: TEMPORAL
                Attribute_Value: _time_value
```

Figure 5: Frame Representation for <buy>

Therefore, we construct relevant queries pertaining to the user's goal, which is used to gather other senses of the goal as well as other goals which are required to achieve the goal. For instance, a typical OMCS query 'buy house', produces binary predicate structures:

1. (EventForGoalEvent "buy house" "apply for mortgage")
2. (EventForGoalEvent "buy house" "ask for loan")
3. (EventForGoalEvent "buy house" "avoid house with termite")
4. (EventForGoalEvent "buy house" "be careful")
5. (EventForGoalEvent "buy house" "contact real estate agent")

6 (EventForGoalEvent "buy house" "contact your local real estate agent")

Moving Into the Expert Domain

At this point we do some handholding with the user to better define the goal. We split this phase into three parts.

First, we use the concepts as queries and crawl the web to get the most relevant links that offer information about the goal. So, from the earlier example of "buy house", common sense comes back with facts like "real estate". The links that CSI returns will pertain to contacting real-estate agents and buying a house. The web provides a wealth of well-conducted research on various topics hence offering the expertise required to narrow down the goal [18]. We carefully extract the best links and display it in a menu along with an in-built browser for the user to navigate.

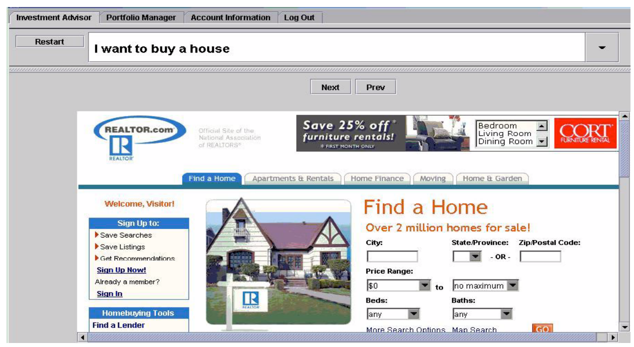

Figure 6: An embedded WWW browser functionality

Second, the user now navigates the web to get more information about the goal. While this is happening, our tool is "listening" to the hyperlinks (Figure 6). When the user finally closes on a price or value it is passed to the system and in the backend the current URL is captured for two reasons. One is to be able to return to the site at a later point either for debugging purposes or to redo the selection. Two is to extract other options that the tool can suggest to the user, if the current choice was not good.

Part three of our expert system is an information filter where an agent goes out to the web looking for financial information particularly pertaining to making investments. The search is intelligent in that it looks at different industries and companies and extracts the factors that affect the performance of the market, like the volatility, price-earnings ratio, etc. At the point this filter is triggered, the agent has at its disposal the asset-allocation determined earlier. So the input to the filter will be the expected performance of the various investments (like stocks, bonds, money market) in order to meet the user's goal and if necessary make a profit.

Common Sense Investment Strategy

The initial asset allocation is determined from the users' goal, timeline and risk tolerance. Now, we delve into each allocation and use a combination of common-sense [19] and expert knowledge to pick the top performing industries and companies the user may consider investing in and we explain the reasons behind making this selection (Figure 7).

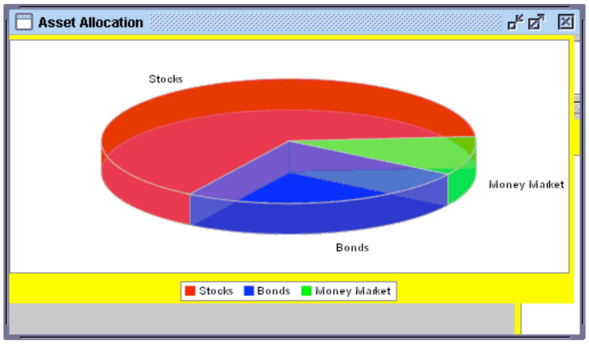

Figure 7: A Sample Asset Allocation

The typical domain specific common sense facts used are as follows:

1. ‘high risk’ -> ‘high return’
2. ‘high return’ -> ‘invest in stocks’
3. (PropertyOf "diversified stock" "good growth with high consistency over long term")
4. (PropertyOf "good stock" "larger the growth rate of dividends and earnings")
5 (CapableOf "high stock allocation" "good return for small amount of capital")

A typical explanation of the kind of common sense that gets used in picking a stock is as follows (Figure 8):

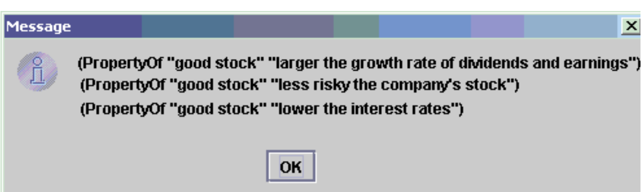

Figure 8: Financial Common Sense predicates

Future Directions

We are working towards gathering more investment-related common sense knowledge. We would also like to do better usability and risk analysis and add functionality to the expert financial engine. For now, we have limited our allocation to just stocks, bonds and money market. This may be extended to retirement funds etc.

ACKNOWLEDGEMENTS

We would like to thank Hugo Liu and Push Singh for numerous suggestions.

REFERENCES

1. http: //www.vanguard.com
2. http://www.fool.com
3. http://www.nasdaq.com
4. http://www.morningstar.com
5. http://www.investopedia.com
6. http://finance.yahoo.com
7. http://www.cnn.com/money
8. Henry Lieberman and Earl Wagner, End-User Debugging for Electronic Commerce, ACM Conference on Intelligent User Interfaces, Miami Beach, January 2003.
9. Personal Choice Point: Helping Users Visualize What It Means to Buy a BMW, Proceedings of the 2003 International Conference on Intelligent User Interfaces, January 2003, Miami, Florida
10. Jim Blythe, Jihie Kim, Surya Ramchandran, and Yolanda Gil, An Integrated Environment for Knowledge Acquistion, Proceeding of the 2001 International conference on Intelligent User Interfaces (IUI-2001), Santa Fe, New Mexico, January 2001.
11. Hugo Liu, Henry Lieberman, Ted Selker. (2002). GOOSE: A Goal-Oriented Search Engine with Commonsense. In De Bra, Brusilovsky, Conejo (Eds.): Adaptive Hypermedia and Adaptive Web-Based Systems, Second International Conference, AH 2002, Malaga, Spain, May 29-31, 2002, Proceedings. Lecture Notes in Computer Science 2347 Springer 2002, ISBN 3-540-43737-1, pp. 253-263.
12. Fillmore, Charles J. (1968): The case for case. In Bach and Harms (Ed.): Universals in Linguistic Theory (pp. 1-88), Holt, Rinehart, and Winston, New York.
13. Fillmore, Charles J. (1976): Frame semantics and the nature of language. In Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech, Volume 280 (pp. 20-32).
14. Fillmore, Charles J. (1982): Frame semantics. In Linguistics in the Morning Calm (pp. 111-137), Hanshin Publishing Co., Seoul, South Korea.
15. Fillmore, Charles J. (1985): Frames and the semantics of understanding. In Quaderni di Semantica, Vol 6, No. 2 (pp. 222-254).
16. Timothy Chklovski. LEARNER: A System for Acquiring Commonsense Knowledge by Analogy. In Proceedings of Second International Conference on Knowledge Capture (K-CAP 2003). October 2003.
17. The Society of Mind, Marvin Minsky.
18. Hugo Liu and Henry Lieberman (2002). Robust photo retrieval using world semantics. In Proceedings of the LREC 2002 Workshop on Creating and Using Semantics for Information Retrieval and Filtering: State-of-the-art and Future Research, Las Palmas, Canary Islands, pp. 15-20.
19. Henry Lieberman and Barbara Barry, Collecting commonsense experiences. 2003. Proceedings of the Second International Conference on Knowledge Capture (K-CAP 2003). Florida, USA.